In the claims

1. (previously presented) A cache-coherent system comprising:

a memory having a plurality of memory units;

a plurality of nodes employing a coherence protocol to maintain cache coherence of the memory;

a cache within each node to temporarily store contents of the plurality of memory units; and,

logic within each node to determine whether a cache miss relating to a memory unit should be transmitted only to a sub-plurality of nodes lesser in number than the plurality of nodes but greater than one, based on a criteria.

2. (original) The system of claim 1, wherein the criteria includes whether, to ultimately reach an owning node for the memory unit, such transmission is likely to reduce total communication traffic among the plurality of nodes and unlikely to increase latency as compared to broadcasting the cache miss to all of the plurality of nodes.

3. (previously presented) The system of claim 1, wherein the logic within each node is to determine whether the node is a home node for the memory unit to which the cache miss relates in determining that transmission to the sub-plurality of nodes lesser in number than the plurality of nodes is likely to reduce total communication traffic among the plurality of nodes and unlikely to increase latency to ultimately reach the owning node for the memory unit.

4. (previously presented) The system of claim 3, wherein the sub-plurality of nodes comprises an owning node for the memory unit as stored at a directory of the home node.

5. (previously presented) The system of claim 1, wherein the logic within each node is to determine whether the cache of the node has stored a hint as to a potential owning node for the memory unit as a result of an earlier event in determining that transmission to the sub-plurality of nodes lesser in number than the plurality of nodes is likely to reduce total communication traffic among the plurality of nodes and unlikely to increase latency to ultimately reach the owning node for the memory unit.

6. (original) The system of claim 5, wherein the event includes an invalidation of the memory unit by the potential owning node.

7. (previously presented) The system of claim 5, wherein the sub-plurality of nodes comprises a home node of the memory unit and the potential owning node for the memory unit.

8. (original) The system of claim 1, wherein the logic within each node is to determine whether the memory unit relates to a predetermined memory sharing pattern encompassing the one or more nodes in determining that transmission to the one or more nodes lesser in number than the plurality of nodes is likely to reduce total communication traffic among the plurality of nodes and unlikely to increase latency to ultimately reach the owning node for the memory unit.

9. (previously presented) A method comprising:

determining at a first node whether a cache miss relating to a memory unit of a shared memory system of a plurality of nodes including the first node and employing a coherence protocol should be selectively broadcast only to a sub-plurality of nodes lesser in number than the plurality of nodes but greater than one based on a criteria;

in response to determining that the cache miss should be selectively broadcast only to the sub-plurality of nodes, selectively broadcasting the cache miss by the first node only to the sub-plurality of nodes.

10. (previously presented) The method of claim 9, further comprising, in response to determining that the cache miss should not be selectively broadcast to the sub-plurality of nodes, broadcasting the cache miss by the first node to all of the plurality of nodes.

11. (original) The method of claim 9, wherein the criteria includes whether selective broadcasting is likely to reduce total communication traffic among the plurality of nodes and unlikely to increase latency as compared to just broadcasting the cache miss to all of the plurality of nodes to reach an owning node for the memory unit.

12. (previously presented) The method of claim 9, wherein determining whether the cache miss should be selectively broadcast to the sub-plurality of nodes comprises determining whether the first node is a home node for the memory unit, such that selectively broadcasting the cache miss to the sub-plurality of nodes comprises selectively broadcasting the cache miss to one node of the plurality of nodes as an owning node for the memory unit as stored at a directory of the first node as the home node for the memory unit.

13. (previously presented) The method of claim 9, wherein determining whether the cache miss should be selectively broadcast to the sub-plurality of nodes comprises determining whether the first node has a pre-stored hint as to a potential owning node for the memory unit, such that selectively broadcasting the cache miss to the sub-plurality of nodes comprises selectively broadcasting the cache miss both to a home node of the memory unit and to the potential owning node for the memory unit.

14. (previously presented) The method of claim 9, wherein determining whether the cache miss should be selectively broadcast to the sub-plurality nodes comprises determining whether the memory unit relates to a predetermined memory sharing pattern encompassing the sub-plurality of nodes, such that selectively broadcasting the cache miss to the sub-plurality of nodes comprises selectively broadcasting the cache miss to the sub-plurality of nodes.

15. (previously presented) A method comprising:

determining at a first node whether a cache miss relating to a memory unit of a shared memory system of a plurality of nodes including the first node and employing a coherence protocol should be selectively broadcast only to sub-plurality of nodes lesser in number than the plurality of nodes but greater than one, based on whether selective broadcasting is likely to reduce total communication traffic among the plurality of nodes and unlikely to increase latency as compared to just broadcasting the cache miss to all of the plurality of nodes to reach an owning node for the memory unit; and,

in response to determining that the cache miss should be selectively broadcast only to the sub-plurality of nodes, selectively broadcasting the cache miss by the first node only to the sub-plurality of nodes.

16. (previously presented) A method comprising:

determining at a first node whether a cache miss relating to a memory unit of a shared memory system of a plurality of nodes including the first node should be selectively broadcast only to a sub-plurality of nodes of the plurality of nodes, based on whether the first node is a home node for the memory unit;

in response to determining that the cache miss should be selectively broadcast only to the sub-plurality of nodes, selectively broadcasting the cache miss by the first node only to the sub-plurality of nodes;

otherwise, determining at the first node whether the memory unit relates to a predetermined memory sharing pattern encompassing a sub-plurality of the plurality of nodes smaller in number than the plurality of nodes; and,

in response to determining that the memory unit relates to the predetermined memory sharing pattern, selectively broadcasting the cache miss by the first node to the sub-plurality of the plurality of nodes.

17. (previously presented) A node of a system having a plurality of nodes comprising:

local memory for which the node is a home node and that is shared among the plurality of nodes;

a directory to track which of the plurality of nodes has cached or modified the local memory of the node;

a cache to temporarily store contents of the local memory and memories of other ones of the plurality of nodes; and,

logic to determine whether a cache miss relating to a local memory should be transmitted only to a sub-plurality of nodes lesser in number than the plurality of nodes but greater than one based on whether, to ultimately reach an owning node for the local memory, such transmission is likely to reduce total communication traffic among the plurality of nodes and unlikely to increase latency as compared to broadcasting the cache miss to all of the plurality of nodes.

18.-21. (cancelled)